

CLAIMS

I claim:

1 1. A method in a data processing system for identifying a subnet address
2 range for a subnet being used in a network, the network utilizing addresses each comprising
3 an ordered series of a fixed number of bits, comprising:

4 (a) determining the addresses of a plurality of hosts within the network;

5 (b) identifying an address prefix for which at least one determined host
6 address has the form "<identified address prefix> 0 <any address suffix>", and for which at
7 least one determined address has the form "<identified address prefix> 1 <any address
8 suffix>";

9 (c) determining whether the plurality of determined host addresses include
10 either the address "<identified address prefix> 0 <address suffix comprised entirely of 1s>"
11 or the address "<identified address prefix> 1 <address suffix comprised entirely of 0s>";

12 (d) if the plurality of determined host addresses include either the address
13 "<identified address prefix> 0 <address suffix comprised entirely of 1s>" or the address
14 "<identified address prefix> 1 <address suffix comprised entirely of 0s>", identifying the
15 range of addresses that all begin with the identified address prefix as a subnet address range
16 for a subnet being used in a network;

17 (e) if the plurality of determined host addresses include neither the address
18 "<identified address prefix> 0 <address suffix comprised entirely of 1s>" nor the address
19 "<identified address prefix> 1 <address suffix comprised entirely of 0s>":

20 . (1) sending one or more test packets requesting a reply each to a
21 determined host address having the form "<identified address prefix> 0 <any address
22 suffix>" from an address having the form "<identified address prefix> 1 <any address
23 suffix>";

24 (2) sending one or more test packets requesting a reply each to a
25 determined host address having the form "<identified address prefix> 1 <any address
26 suffix>" from an address having the form "<identified address prefix> 0 <any address
27 suffix>";

28 (3) determining the number of test packets sent where a reply to the
29 test packet was sent directly from the test packet's destination address to the test packet's
30 source address; and

1 2. The method of claim 1, further comprising:

(A) repeating steps (a)-(e)(5)(B) for address prefixes beginning with "<identified address prefix> 0"; and

(B) repeating steps (a)-(e)(5)(B) for address prefixes beginning with “<identified address prefix> 1”.

1 3. The method of claim 1, further comprising:

generating a subnet tree data structure containing leaf nodes each representing
a different determined host address, each leaf node having an ordered series of ancestor
nodes that each represent an address prefix that is a prefix of the determined host address, the
length of the address prefixes represented by the ancestor nodes decreasing as distance from
the leaf node in the series of the ancestor nodes increases; and

7 traversing the generated subnet tree to identify the address prefix.

1 4. The method of claim 1 wherein the network is an Internet Protocol
2 network, and wherein the addresses are Internet Protocol addresses.

1 5. A method in a data processing system for identifying subnet address
2 ranges for a subnets being used in a network, comprising:

3 determining a plurality of addresses of hosts in the network;

4 accessing a binary tree, the binary tree having a root node having no parents,
5 parent nodes including the root node each having two child nodes, and leaf nodes having no
6 children, such that the root node represents the entire range of addresses available in the
7 network, such that each child node in a pair of child nodes represents a distinct half of the
8 range represented by the parent node of the pair of child nodes, and such that each leaf node
9 represents a single network address that is within the address ranges represented by all of the
10 ancestors of the leaf node, each determined host address being represented by a leaf node;

11 traversing the binary tree in preorder to identify candidate nodes, both child
12 nodes of each candidate node having one or more descendant leaf nodes representing a
13 determined host address;

14 testing the address range represented by each visited candidate node to
15 determine whether the address range is a subnet address range for a subnet being used in a
16 network;

17 if testing indicates that a visited candidate node represents an address range
18 that is a subnet address range for a subnet being used in a network, identifying the visited
19 candidate node as a subnet node; and

20 skipping, in the traversal, any candidate nodes that are descendants of a
21 identified subnet node.

1 6. The method of claim 5 wherein testing comprises, for the two subranges
2 represented by the child nodes of the candidate node:

3 sending one or more packets each from a source address to a destination
4 address, each packet requesting a reply, the source and destination addresses being in
5 different subranges for each packet;

6 for each packet, determining whether a reply to the packet is sent directly from
7 the destination address back to the source address; and

8 if, for a number of packets exceeding a threshold number, a reply to the packet
9 is sent directly from the destination address back to the source address, determining that the
10 candidate node represents an address range that is a subnet address range for a subnet being
11 used in a network.

1 7. The method of claim 5 wherein testing comprises, for the two subranges
2 represented by the child nodes of the candidate node:

3 selecting the address within each subrange that is closest to the addresses of the
4 other subrange;

5 determining whether the network contains a host responding to either of the
6 selected addresses; and

7 if the network contains a host responding to either of the selected addresses,
8 determining that the candidate node represents an address range that is a subnet address range
9 for a subnet being used in a network.

10 8. The method of claim 5, further comprising, before traversing the binary
11 tree, trimming the binary tree by deleting nodes not on the path between the root node and
any leaf node representing a determined host address.

1 9. The method of claim 5, further comprising, before traversing the binary
2 tree, trimming the binary tree by deleting all nodes not on a path between the root node and
3 any leaf node representing a determined host address.

1 10. A computer-readable medium whose contents cause a data processing
2 system to identify subnet address ranges for a subnets being used in a network by:

3 receiving a plurality of addresses of hosts in the network;

4 accessing a binary tree, the binary tree having a root node having no parents,
5 parent nodes including the root node each having two child nodes, and leaf nodes having no
6 children, such that the root node represents the entire range of addresses available in the
7 network, such that each child node in a pair of child nodes represents a distinct half of the
8 range represented by the parent node of the pair of child nodes, and such that each leaf node
9 represents a single network address that is within the address ranges represented by all of the
10 ancestors of the leaf node, each received host address being represented by a leaf node;

11 traversing the binary tree in preorder to identify candidate nodes, both child
12 nodes of each candidate node having one or more descendant leaf nodes representing a
13 received host address;

14 testing the address range represented by each candidate node in the traversal
15 visited to determine whether the address range is a subnet address range for a subnet being
16 used in a network;

17 if testing indicates that a visited candidate node represents an address range
18 that is a subnet address range for a subnet being used in a network, identifying the visited
19 candidate node as a subnet node; and

20 skipping, in the traversal, any candidate nodes that are descendants of a
21 identified subnet node.

1 11. The computer-readable medium of claim 10 wherein testing comprises,
2 for the two subranges represented by the child nodes of the candidate node:

3 sending one or more packets each from a source address to a destination
4 address, each packet requesting a reply, the source and destination addresses being in
5 different subranges for each packet;

6 for each packet, determining whether a reply to the packet is sent directly from
7 the destination address back to the source address; and

8 if, for a number of packets exceeding a threshold number, a reply to the packet
9 is sent directly from the destination address back to the source address, determining that the
10 candidate node represents an address range that is a subnet address range for a subnet being
11 used in a network.

1 12. The computer-readable medium of claim 10 wherein testing comprises,
2 for the two subranges represented by the child nodes of the candidate node:

3 selecting the address within each subrange that is closest to the addresses of the
4 other subrange;

5 determining whether the network contains a host responding to either of the
6 selected addresses; and

7 if the network contains a host responding to either of the selected addresses,
8 determining that the candidate node represents an address range that is a subnet address range
9 for a subnet being used in a network.

1 13. A method in a data processing system for determining whether a
2 selected range of addresses in a network is entirely within a subnet, comprising:
3 identifying two subranges within the range;
4 sending one or more packets each from a source address to a destination
5 address, each packet requesting a reply, the source and destination addresses being in
6 different subranges for each packet;
7 for each packet, determining whether a reply to the packet is sent directly from
8 the destination address back to the source address; and
9 if, for a number of packets exceeding a threshold number, a reply to the packet
10 is sent directly from the destination address back to the source address, determining that the
11 range of addresses is entirely within a subnet.

1 14. The method of claim 13 wherein the identifying identifies contiguous,
2 mutually exclusive subranges each containing half as many addresses as the range.

1 15. A computer-readable medium whose contents cause a data processing
2 system to determining whether a selected range of addresses in a network is entirely within a
3 subnet by:
4 dividing the range into two subranges;
5 sending one or more packets each from a source address to a destination
6 address, each packet requesting a reply, the source and destination addresses being in
7 different subranges for each packet;
8 for each packet, determining whether a reply to the packet is sent directly from
9 the destination address back to the source address; and
10 if, for a number of packets exceeding a threshold number, a reply to the packet
11 is sent directly from the destination address back to the source address, determining that the
12 range of addresses is entirely within a subnet.

1 16. A data processing system for determining whether a selected range of
2 addresses in a network is entirely within a subnet, comprising:
3 a range identification subsystem that identifies two subranges within the range;

4 a packet transmission subsystem that sends one or more packets each from a
5 source address to a destination address, each packet requesting a reply, the source and
6 destination addresses being in different subranges for each packet;

7 a reply monitoring subsystem that, for each packet, determines whether a reply
8 to the packet is sent directly from the destination address back to the source address; and

9 a result determination subsystem that determines that the range of addresses is
10 entirely within a subnet if, for a number of packets exceeding a threshold number, a reply to
11 the packet is sent directly from the destination address back to the source address.

1 17. A method in a data processing system for determining whether a
2 selected range of addresses in a network is entirely within a subnet, comprising:

3 dividing the range into two subranges;

4 within each subrange, selecting one address;

5 determining whether the network contains a device responding to either of the
6 selected addresses; and

7 if the network contains a device responding to either of the selected addresses,
8 determining that the range of addresses is entirely within a subnet.

1 18. The method of claim 17 wherein the identifying identifies contiguous,
2 mutually exclusive subranges each containing half as many addresses as the range, and
3 wherein the selecting selects the address in each subrange that is closest to the addresses of
4 the other subrange.

1 19. A computer-readable medium whose contents cause a data processing
2 system to determine whether a selected range of addresses in a network is entirely within a
3 subnet, comprising:

4 identifying two subranges within the range;

5 within each subrange, selecting one address;

6 determining whether the network contains a device responding to either of the
7 selected addresses; and

8 if the network contains a device responding to either of the selected addresses,
9 determining that the range of addresses is entirely within a subnet.

1 20. A data processing system for determining whether a selected range of
2 addresses in a network is entirely within a subnet, comprising:

3 a range identification subsystem that identifies two subranges within the range;
4 an address selection system that selects one address within each subrange;
5 a network device sensing subsystem that determines whether the network
6 contains a device responding to either of the selected addresses; and
7 a result determination subsystem that determines that the range of addresses is
8 entirely within a subnet if the network contains a device responding to either of the selected
9 addresses.

1 21. A computer memory containing a subnet identification tree data
2 structure for use in identifying a subnet in a network, a range of addresses being available in
3 the network, some of the addresses available in the network being host addresses, the a
4 subnet identification tree data structure comprising:

5 a root node having two child nodes, the root node representing the entire range
6 of addresses available in the network;

7 a plurality of intermediate nodes each having one parent node and two child
8 nodes, the root node and the intermediate nodes each having two child nodes, each child
9 node in a pair of child nodes representing a distinct half of the range represented by the
10 parent node of the pair of child nodes; and

11 a plurality of leaf nodes having no children, each leaf node representing a
12 single network address that is within the address ranges represented by all of the ancestors of
13 the leaf node, each host address being represented by a leaf node,

14 so that the subnet identification tree data structure may be traversed to identify
15 an intermediate node representing an address range corresponding to a subnet in the network.

1 22. The computer memory of claim 21 wherein the [a] subnet identification
2 tree data structure contains only nodes that are on a path from the root node to one of the leaf
3 nodes representing a host address.